

7 Claims

What is claimed in the present invention is

1. A system for storing one or more time series comprising:

- (a) a language for describing the storing of the one or more time series; and
- (b) a subsystem storing the one or more time series in accordance with said language.

2. A system for storing one or more time series as in claim 1 wherein said language comprises one or more attributes for representing one or more fields of the time series.

3. A system for storing one or more time series as in claim 2 wherein said subsystem comprises one or more rules for describing the storing of said one or more fields in accordance with said one or more attributes.

4. A system for storing one or more time series as in claim 3 wherein said subsystem further comprises:

- (a) at least one file name; and
- (b) at least one data file.

5. A system for storing one or more time series as in claim 4 wherein said one or more rules determine whether the fields are stored in said at least one file name or said at least one data file depending on values of said one or more attributes.

6. A system for storing one or more time series as in claim 1 wherein said language comprises one or more members of the set consisting of a leaf node, a non-leaf node, a type and a hint for describing one or more fields from the time series.

7. A system for storing one or more time series as in claim 6 wherein said subsystem comprises:

- (a) at least one file name; and
- (b) at least one data file.

8. A system for storing one or more time series as in claim 7 wherein said subsystem further comprises one or more rules for determining how to store at least one of the fields from the time series.

9. A system for storing one or more time series as in claim 8 wherein said rules comprise:

- (a) storing the at least one field in said file name if the field is said non-leaf node.

10. A system for storing one or more time series as in claim 8 wherein said rules comprise: storing the at least one field in said filename if:

- (a) the field is said leaf node; and

(b) the field has a fixed value for said hint.

11. A system for storing one or more time series as in claim 8 wherein said rules comprise: storing the at least one field in said data file if:

(a) the field is said leaf node;

(b) the field has a variable value for said hint; and

(c) the field has a constant size.

12. A system for storing one or more time series as in claim 8 wherein said rule comprise: storing a universal matching symbol in the filename if:

(a) the at least one field is said leaf node;

(b) the field has a variable value for said hint; and

(c) the field has a constant size.

13. A system for storing one or more time series as in claim 8 wherein said subsystem further comprises at least one secondary storage.

14. A system for storing one or more time series as in claim 13 wherein said rules comprise: store the at least one field in the secondary storage if:

(a) the field is said leaf node;

(b) the field has a variable value for said hint; and

(c) the field has a constant size.

15. A system for storing one or more time series as in claim 13 wherein said rules comprise: store a universal matching symbol in the filename if:

(a) the field is said leaf node;

(b) the field has a variable value for said hint; and

(c) the field has a constant size.

16. A system for storing one or more time series as in claim 13 wherein said rule comprises: copy an offset of the secondary storage into said data file if:

(a) the field is said leaf node;

(b) the field has a variable value for said hint; and

(c) the field has a constant size.

17. A system for storing one or more time series as in claim 1 wherein said language defines an exchange rate of one or more currencies.

18. A system for storing one or more time series as in claim 1 wherein said language defines a deposit of a currency for a period.

19. A system for storing one or more time series as in claim 1 wherein said language defines a quote.
20. A system for storing one or more time series as in claim 19 wherein said quote comprises one or more number of the set consisting of a bid, an ask, a bank and a source.
- 5 21. A system for storing one or more time series as in claim 1 wherein said language comprises a transaction.
22. A system for storing one or more time series as in claim 21 wherein said transaction comprises an exchange of a currency from a seller to a buyer.
- 10 23. A system for storing one or more time series as in claim 22 wherein said transaction further comprises one or more members of the set consisting of a price, a volume and a source.
24. A system for storing one or more time series as in claim 1 wherein said language comprises one or more regular expressions.
- 15 25. A system for storing one or more time series as in claim 1 wherein said language comprises one or more statements for defining one or more ticks in the time series.
26. A system for storing one or more time series as in claim 1 wherein said language is recursive.
- 20 27. A system for managing one or more time series comprising:
(a) a language defining a first one of the time series as a subset of a second one of the time series.
28. A system for managing one or more time series as in claim 27 wherein said second time series is a universal time series representing all recordable events.
29. A system for retrieving desired data from one or more time series comprising:
(a) at least one request comprising one or more restrictions for defining the desired data; and
25 (b) at least one utility retrieving data from the one or more time series that satisfies said one or more restrictions.
30. A system for retrieving desired data from one or more time series as in claim 29 further comprising:
30 (a) one or more rules for selecting one or more files comprising the one or more time series that satisfy said one or more restrictions.
31. A system for retrieving desired data from one or more time series as in claim 30 further comprising a language, said language defining one or more attributes for said one or more restrictions.

32. A system for retrieving desired data from one or more time series as in claim 31 wherein said one or more attributes comprise one or more members of the set consisting of a node type, a data type and a hint.

33. A system for retrieving desired data from one or more time series as in claim 32 wherein said one or more rules comprise: for at least one of said restrictions,

(a) select said one or more filename that match said at least one restriction if said node type of said at least one restriction is a non-leaf.

34. A system for retrieving desired data from one or more time series as in claim 32 wherein said one or more rules comprise: for at least one of said restrictions,

(a) select said one or more filenames having a universal matching symbol corresponding to said at least one restriction if:

(b) said node type of said at least one restriction is a leaf; and

(c) said hint of said restriction is a variable.

35. A system for retrieving desired data from one or more time series as in claim 33 wherein said one or more rules comprise: for at least one of said restrictions,

(a) select said one or more filenames that match said at least one restriction if:

(b) said node type of said at least one restriction is a leaf; and

(c) said hint of said at least one restriction is fixed.

36. A system for retrieving desired data from one or more time series as in claim 29 further comprising:

(a) at least one cursor for selecting data in the one or more time series that satisfies said one or more restrictions.

37. A system for retrieving desired data from one or more time series as in claim 36 wherein said one or more restrictions comprise one or more members of the set consisting of a base time and a time range.

38. A system for retrieving desired data from one or more time series as in claim 37 wherein said time range specifies the number of data items before said base time.

39. A system for retrieving desired data from one or more time series as in claim 37 wherein said time range specifies the number of data item after said base time.

40. A system for retrieving desired data from one or more time series as in claim 36 wherein said cursor comprises a first method for retrieving the data item after a current time in the one or more time series that satisfies said one or more restrictions.

41. A system for retrieving desired data from one or more time series as in claim 36 wherein said cursor comprises:

(a) a second method for retrieving the data item before a current time in the one or more time series that satisfies said one or more restriction.

42. A system for retrieving desired data from one or more time series as in claim 29 further comprising:

5 (a) a parser for determining said one or more restrictions from said at least one request.

43. A system for retrieving desired data from one or more time series as in claim 29 wherein said one or more restrictions comprise an expression.

44. A system for processing data from one or more time series comprising:

(a) one or more processing modules for processing the data;

10 (b) one or more connections for linking said modules in a network; and

(c) a first subsystem for activating said one or more processing modules and for moving the data through the network.

45. A system for processing data from one or more time series as in claim 44 further comprising a type system comprising:

15 (a) one or more types; and

(b) a relation among said one or more types.

46. A system for processing data from one or more time series as in claim 45 further comprising a grammar to describe said types in said type system.

47. A system for processing data from one or more time series as in claim 45 wherein said one or more processing modules comprise one or more ports.

20

48. A system for processing data from one or more time series as in claim 47 further comprising one or more binding operators for creating said one or more connections to link two or more of said ports.

49. A system for processing data from one or more time series as in claim 48 wherein at least one of said types are assigned to at least one of said ports.

25

50. A system for processing data from one or more time series as in claim 49 wherein said one or more processing modules comprise:

(a) a configure method for checking that said types on said ports that are linked by one of said connections are consistent.

30

51. A system for processing data from one or more time series as in claim 44 wherein said processing modules comprise:

(a) a process data method to process the data.

52. A system for processing data from one or more time series as in claim 51 wherein said subsystem executes said process data method.

53. A system for processing data from one or more time series as in claim 44 wherein at least one datum of the data in the time series has at least one time stamp.

5 54. A system for processing data from one or more time series as in claim 53 wherein said subsystem:

(a) orders said at least one datum of the data according to said time stamp; and

(b) provides said ordered at least one datum to said processing modules.

10 55. A system for processing data from one or more time series as in claim 44 wherein said processing modules comprise one or more ports.

56. A system for processing data from one or more time series as in claim 55 wherein said ports comprise one or more input ports and one or more output ports.

57. A system for processing data from one or more time series as in claim 56 wherein said processing modules further comprise:

15 (a) at least one end of data method to indicate that no more data will be provided to said one or more input ports of said processing modules.

58. A system for processing data from one or more time series as in claim 57 wherein said first subsystem executes said end of data method when said subsystem has no more of the data to provide to said processing module.

20 59. A system for processing data from one or more time series as in claim 56 wherein said processing modules input at least one input datum of the data on said input ports, process said at least one input datum to produce at least one output datum, and output said at least one output datum on said output ports.

25 60. A system for processing data from one or more time series as in claim 59 wherein said processing module further comprises a build-up delay method that computes how much time said processing module needs before said processing module can output said at least one output datum that is meaningful.

30 61. A system for processing data from one or more time series as in claim 59 wherein said processing modules further comprise one or more timer methods to process one or more timers.

62. A system for processing data from one or more time series as in claim 61 wherein said one or more timers indicate when said processing modules should output said at least one output datum on said output ports.

63. A system for processing data from one or more time series as in claim 62 wherein said processing modules compute an average of input data and output said average at its said outputs at time intervals.

5

64. A system for processing data from one or more time series as in claim 63 wherein said time intervals are hourly.

65. A system for processing data from one or more time series as in claim 59 wherein said processing module comprise:

(a) at least one end of run method to indicate that said processing module should output any remaining said at least one output datum.

10

66. A system for processing data as in claim 65 wherein said first subsystem executes said end of run method.

67. A system for processing data from one or more time series as in claim 44 wherein said processing modules comprise one or more variables defining a state of said processing modules.

15

68. A system for processing data from one or more time series as in claim 44 wherein each of said processing modules execute independently of others of said processing modules.

69. A system for processing data from one or more time series as in claim 44 wherein said processing module further comprise:

(a) one or more timer methods to process one or more timers.

20

70. A system for processing data from one or more time series as in claim 69 wherein said first subsystem executes said timer methods.

71. A system for processing data from one or more time series as in claim 44 wherein the network is directed acyclic graph.

25

72. A system for processing data from one or more time series as in claim 44 wherein said processing modules comprise one or more members of the set consisting of producer-only modules that output data but do not input data, producer-consumer modules that input data and output data; and consumer-only modules that input data but do not output data.

30

73. A system for processing data from one or more time series as in claim 44 wherein said processing modules comprise one or more members of the set consisting of modules that read data from a repository, modules that perform financial calculations, modules that perform computations, modules that perform statistical analysis, modules that compute histograms, and modules that write data to the repository.

35

74. A system for processing data from one or more time series as in claim 73 wherein the computations comprise one or more members from the set consisting of:

(a) derivatives, volatility, generation of regular time series.

75. A system for processing data from one or more time series as in claim 73 wherein the financial calculation comprise a cross-rate with foreign exchange data.

5 76. A system for processing data from one or more time series as in claim 73 wherein the statistical analysis comprise one or more members of the set consisting of correlation, moving correlation and least square fit.

77. A system for processing data from one or more time series as in claim 73 wherein the histogram comprise one or more members of the set consisting of probability distribution, and conditional average and intra-week average.

10 78. A system for processing data from one or more time series as in claim 44 wherein said processing module process data from one or more time intervals.

79. A system for processing data from one or more time series as in claim 44 further comprising at least one start time and at least one end time.

15 80. A system for processing data from one or more time series as in claim 79 wherein said processing modules begin the processing of the data at said start time and continue processing the data until said end time.

81. A system for processing data from one or more time series as in claim 79 wherein said subsystem passes said start time and said end time to said processing modules.

20 82. A system for processing data from one or more time series as in claim 44 wherein at least one of said processing modules computes and produces a running average of at least one datum from the data that said processing modules receive.

83. A system for processing data from one or more time series as in claim 44 wherein at least one of said processing modules outputs an average of its said at least one received datum at every nth one of its said at least one received datum.